

ROBUSTNESS, CANALYZING FUNCTIONS AND SYSTEMS DESIGN

JOHANNES RAUH AND NIHAT AY


ABSTRACT. We study a notion of robustness of a Markov kernel that describes a system of several input random variables and one output random variable. Robustness requires that the behaviour of the system does not change if one or several of the input variables are knocked out. If the system is required to be robust against too many knockouts, then the output variable cannot distinguish reliably between input states and must be independent of the input. We study how many input states the output variable can distinguish as a function of the required level of robustness.

Gibbs potentials allow a mechanistic description of the behaviour of the system after knockouts. Robustness imposes structural constraints on these potentials. We show that interaction families of Gibbs potentials allow to describe robust systems.

Given a distribution of the input random variables and the Markov kernel describing the system, we obtain a joint probability distribution. Robustness implies a number of conditional independence statements for this joint distribution. The set of all probability distributions corresponding to robust systems can be decomposed into a finite union of components, and we find parametrizations of the components. The decomposition corresponds to a primary decomposition of the conditional independence ideal and can be derived from more general results about generalized binomial edge ideals.

1. INTRODUCTION

Consider a stochastic system of n input nodes and one output node:



As shown in [1], there are two ingredients to robustness:

- (1) If one or several of the input nodes are removed, the system behaviour should not change too much (“small exclusion dependence”).
- (2) A causal contribution of the input nodes on the output nodes.

The second point is strictly necessary: If the behaviour of the output does not depend on the inputs at all, then it is usually not affected by a knockout of a subset of the inputs, but this exclusion independence is trivial.

In this paper we do not use the information theoretic measures proposed in [1]. Instead, we start with a simple model of exclusion independence: We study systems in which the behaviour of the output node does not change when one or more of the

Date: October 30, 2012.

Key words and phrases. robustness, conditional independence, Markov kernels.

input nodes are knocked out. We formalize our robustness requirements in terms of a *robustness specification* \mathcal{R} , which consists of pairs (R, x_R) , where R is a subset of the inputs and x_R is a joint state of the inputs in R . Let \mathcal{S} be a set of possible states of the input nodes. The system is \mathcal{R} -robust in \mathcal{S} , if the behaviour of the system does not change if the inputs not in R are knocked out, provided that the inputs in R are in the state x_R and the current state of all inputs belongs to \mathcal{S} .

If the robustness specification \mathcal{R} is too large, or if the set \mathcal{S} is too large, then in any \mathcal{R} -robust system the output does not depend on the input at all. In general, the behaviour of the system is restricted by robustness requirements. Therefore, to study the causal contribution of the input nodes on the output nodes, we investigate how varied the behaviour of a system can be, given both \mathcal{R} and \mathcal{S} . More precisely, robustness specifications imply that the system cannot distinguish all input states, and we may ask how many states the system can discern. This question is related to the topic of error detecting codes, see Remark 6.

This paper is organized as follows: Section 2 contains our basic setting and definitions. We find several equivalent formulations of our notion of robustness. Moreover, we study the question how many states an \mathcal{R} -robust system can distinguish. Section 3 shows that our definitions generalize the notions of canalyzing [9] and nested canalyzing functions [8], which have been studied before in the context of robustness. Section 4 proposes to model the different behaviours of a system under various knockouts using a family of Gibbs potentials. Robustness implies various constraints on these potentials. Section 4 discusses the probabilistic behaviour of the whole system, including its inputs, when the input variables are distributed to some fixed input distribution. The set of all joint probability distributions is found such that the system is \mathcal{R} -robust for all input states with non-vanishing probability.

Some of our results in Section 5 can also be derived from recent algebraic results in [13] about generalized binomial ideals. These ideals generalize the binomial edge ideals of [6] and [12]. Similar ideals have recently been studied in the paper [14], which discusses what we call $(n-1)$ -robustness in Section 6. In this paper we give self-contained proofs that are also accessible to readers not acquainted to the language of commutative algebra. We comment on the relation to the algebraic results in Remark 25.

2. ROBUSTNESS AND CANALYZING FUNCTIONS

We consider n input nodes, denoted by $1, 2, \dots, n$, and one output node, denoted by 0 . For each $i = 0, 1, \dots, n$ the state of node i is a discrete random variable X_i taking values in the finite set \mathcal{X}_i of cardinality d_i . The input state space is the set $\mathcal{X}_{\text{in}} = \mathcal{X}_1 \times \dots \times \mathcal{X}_n$, and the joint state space is $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_{\text{in}}$. For any subset $S \subseteq \{0, \dots, n\}$ write X_S for the random vector $(X_i)_{i \in S}$; then X_S is a random variable with values in $\mathcal{X}_S = \prod_{i \in S} \mathcal{X}_i$. For $S = [n] := \{1, \dots, n\}$ we also write X_{in} instead of $X_{[n]}$. For any $x \in \mathcal{X}$, the *restriction* of x to a subset $S \subseteq \{0, \dots, n\}$ is the vector $x|_S \in \mathcal{X}_S$ with $(x|_S)_i = x_i$ for all $i \in S$. In contrast, the notation x_S will refer to an arbitrary element of \mathcal{X}_S .

As a model for the computation of the output from the input, we use a stochastic map (Markov kernel) κ from \mathcal{X}_{in} to \mathcal{X}_0 , that is, κ is a function that assigns to each $x \in \mathcal{X}_{\text{in}}$ a probability distribution $\kappa(x)$ for the output X_0 . Such a stochastic map κ can be represented by a matrix, with matrix elements $\kappa(x; x_0)$, $x \in \mathcal{X}_{\text{in}}$, $x_0 \in \mathcal{X}_0$, satisfying $\sum_{x_0 \in \mathcal{X}_0} \kappa(x; x_0) = 1$ for all $x \in \mathcal{X}_{\text{in}}$. For each $x \in \mathcal{X}_{\text{in}}$ the probability

distribution $\kappa(x)$ models the behaviour of X_0 when the input variables are in the state x . When the input is distributed according to some input distribution p_{in} , then the joint distribution p of input and output variables satisfies

$$p(X_0 = x_0, X_{\text{in}} = x) = p_{\text{in}}(X_{\text{in}} = x)\kappa(x; x_0).$$

If $p_{\text{in}}(X_{\text{in}} = x) > 0$, then $\kappa(x)$ can be computed from the joint probability distribution p and equals the conditional distribution of X_0 , given that $X_{\text{in}} = x$.

When a subset S of the input nodes is knocked out and only the nodes in $R = [n] \setminus S$ remain, then the behaviour of the system changes. Without further assumptions, the post-knockout function is not determined by κ and has to be specified separately. We model the post-knockout function by a further stochastic map $\kappa_R : \mathcal{X}_R \times \mathcal{X}_0 \rightarrow [0, 1]$. A complete specification of the system is given by the family $(\kappa_A)_{A \subseteq [n]}$ of all possible post-knockout functions, which we refer to as *functional modalities*. As a shorthand notation we denote functional modalities by (κ_A) . The stochastic map κ itself, which describes the normal behaviour of the system without knockouts, can be identified with $\kappa_{[n]}$.

What does it mean for functional modalities to be robust? Assume that the input is in state x , and that we knock out a set S of inputs. Denoting the remaining set of inputs by R , we say that (κ_A) is robust in x against knockout of S , if $\kappa(x) = \kappa_R(x|_R)$, that is, if

$$(1) \quad \kappa(x; x_0) = \kappa_R(x|_R; x_0) \quad \text{for all } x_0 \in \mathcal{X}_0.$$

Let \mathcal{R} be a collection of pairs (R, x_R) , where $R \subseteq [n]$ and $x_R \in \mathcal{X}_R$. We call such a collection a *robustness specification* in the following. We say that (κ_A) is \mathcal{R} -robust in a set $\mathcal{S} \subseteq \mathcal{X}_{\text{in}}$ if

$$(2) \quad \kappa(x) = \kappa_R(x|_R), \quad \text{whenever } x \in \mathcal{S} \text{ and } (R, x|_R) \in \mathcal{R}.$$

The main example in this section will be the robustness structures

$$\mathcal{R}_k := \{(R, x_R) : R \subseteq [n], |R| \geq k, x_R \in \mathcal{X}_R\}.$$

Equation (1) only compares the functional modality κ_R after knockout with the stochastic map κ that describes the regular behaviour of the unperturbed system. In particular, for $R \subsetneq R' \subseteq [n]$, the functional modality $\kappa_{R'}$ is in no way restricted by (1). Therefore, it may happen that a system that is not robust against a knockout of a set $S' = [n] \setminus R'$ recovers its regular behaviour if we knockout even more nodes. However, this is not the typical situation. Therefore, it is natural to assume that the following holds: If $(R, x_R) \in \mathcal{R}$ and if $R \subsetneq R' \subseteq [n]$, then $(R', x_{R'}) \in \mathcal{R}$ for all $x_{R'} \in \mathcal{X}_{R'}$ with $x_{R'}|_R = x_R$. In this case we call the robustness specification \mathcal{R} *coherent*. For example the robustness structures \mathcal{R}_k are coherent. The notion of coherence will not play an important role in the following, but it is interesting from a conceptual point of view. It is related to the notion of coherency as used e.g. in [3].

By definition, for robust functional modalities (κ_A) the largest functional modality $\kappa_{[n]}$ determines the smaller ones in the relevant points via (2). This motivates the following definition: A stochastic map κ is called \mathcal{R} -robust in \mathcal{S} , if there exist functional modalities (κ_A) with $\kappa = \kappa_{[n]}$ that are \mathcal{R} -robust in \mathcal{S} . More directly, κ is \mathcal{R} -robust in \mathcal{S} if and only if

$$\kappa(x) = \kappa(y), \quad \text{whenever } x, y \in \mathcal{S}, x|_R = y|_R \text{ and } (R, x|_R) \in \mathcal{R}.$$




FIGURE 1. An illustration of Example 1 with $n = 4$. a) The graph $G_{\mathcal{R}_3}$. b) An induced subgraph $G_{\mathcal{R}_3, \mathcal{S}}$. c) The connected components of $G_{\mathcal{R}_3, \mathcal{S}}$. In fact, in this example both connected components are cylinder sets. d) The induced subgraph $G_{\mathcal{R}_2, \mathcal{S}}$, which is connected.

When studying robustness of a stochastic map κ we may always assume that \mathcal{R} is coherent; for if $x|_R = y|_R$ implies $\kappa(x) = \kappa(y)$, then $x|_{R'} = y|_{R'}$ also implies $\kappa(x) = \kappa(y)$, whenever $R \subseteq R' \subseteq [n]$.

For any subset $R \subseteq [n]$ and $x_R \in \mathcal{X}_R$ let

$$C(R, x_R) := \{x \in \mathcal{X}_{\text{in}} : x|_R = x_R\}.$$

be the corresponding cylinder set. Then κ is \mathcal{R} -robust in \mathcal{S} if and only if $\kappa(x) = \kappa(y)$ for all $x, y \in \mathcal{S} \cap C(R, x_R)$ and $(R, x_R) \in \mathcal{R}$. In other words, the stochastic map κ is constant on $\mathcal{S} \cap C(R, x_R)$ for all $(R, x_R) \in \mathcal{R}$.

The following construction is useful to study robust functional modalities: Given a robustness specification \mathcal{R} , define a graph $G_{\mathcal{R}}$ on \mathcal{X}_{in} by connecting two elements $x, y \in \mathcal{X}_{\text{in}}$ by an edge if there is $(R, x_R) \in \mathcal{R}$ such that $x|_R = y|_R = x_R$. Denote by $G_{\mathcal{R}, \mathcal{S}}$ the subgraph of $G_{\mathcal{R}}$ induced by \mathcal{S} .

Example 1. Assume that $\mathcal{X}_i = \{0, 1\}$ for $i = 1, \dots, n$. Then the input state space $\mathcal{X}_{\text{in}} = \{0, 1\}^n$ can be identified with the vertices of an n -dimensional hypercube. The graph $G_{\mathcal{R}_{n-1}}$ is the edge graph of this hypercube (Fig. 1a)). Cylinder sets correspond to faces of this hypercube. If $R \subset [n]$ has cardinality $n-1$, then the cylinder set $C(\mathcal{R}, x_R)$ is an edge, if R has cardinality $n-2$, then $C(\mathcal{R}, x_R)$ is a two-dimensional face. Fig. 1b) shows an induced subgraph of G_3 for $n = 4$. By comparison, the graph $G_{\mathcal{R}_{n-2}}$ has additional edges corresponding to diagonals in the quadrangles of $G_{\mathcal{R}_{n-1}}$. For example, the set of vertices marked black in Figure 1b) is connected in $G_{\mathcal{R}_{n-2}}$, but not in $G_{\mathcal{R}_{n-1}}$ (Fig. 1d)).

Proposition 2. *The following statements are equivalent for a stochastic map κ :*

- (1) κ is \mathcal{R} -robust in \mathcal{S} .
- (2) κ is constant on $\mathcal{S} \cap C(R, x_R)$ for all $(R, x_R) \in \mathcal{R}$.
- (3) κ is constant on the connected components of $G_{\mathcal{R}, \mathcal{S}}$.
- (4) For any probability distribution p_{in} of \mathcal{X}_{in} with $p_{\text{in}}(\mathcal{S}) = 1$ and for all $(R, x_R) \in \mathcal{R}$, the output X_0 is stochastically independent of $X_{[n] \setminus R}$ given $X_R = x_R$.

Proof. The equivalence (1) \Leftrightarrow (2) was already shown.

(2) \Leftrightarrow (3): Condition (2) says that κ is constant along each edge of $G_{\mathcal{R}, \mathcal{S}}$. By iteration this implies (3). In the other direction, the subgraph of $G_{\mathcal{R}, \mathcal{S}}$ induced by $\mathcal{S} \cap C(R, x_R)$ is connected for all $(R, x_R) \in \mathcal{R}$, and therefore (3) implies (2).

(2) \Rightarrow (4): For any $x \in \mathcal{X}_{\text{in}}$ with $p_{\text{in}}(x) > 0$, the conditional distribution of the output given the input satisfies $p(X_0 = x_0 | X_{\text{in}} = x) = \kappa(x; x_0)$. By (2), $\kappa(x; x_0)$ is constant on $C(\mathcal{R}, x|_R) \cap \mathcal{S}$. Hence the conditional distribution does not depend on $X_{[n] \setminus R}$, and so $p(X_0 = x_0 | X_{\text{in}} = x) = p(X_0 = x_0 | X_R = x|_R)$.


(4) \Rightarrow (2): Let p_{in} be the uniform distribution on \mathcal{S} (or any other probability distribution with support \mathcal{S}), and fix $(\mathcal{R}, x_R) \in \mathcal{R}$. By assumption, for any $x \in \mathcal{S}$ with $x|_R = x_R$, the conditional distribution $p(X_0 = x_0 | X_{\text{in}} = x) = \kappa(x; x_0)$ does not depend on $x|_{[n] \setminus \mathcal{S}}$. Therefore, $\kappa(x)$ is constant on $\mathcal{S} \cap C(\mathcal{R}; x_R)$. \square

The choice of the set \mathcal{S} is important: On one hand \mathcal{S} should be large, because otherwise the notion of robustness is very weak. However, if \mathcal{S} is too large, then the equations (1) imply that the output X_0 is (unconditionally) independent of all inputs. Proposition 2 gives a hint how to choose the set \mathcal{S} : The goal is to have as many connected components as possible in $G_{\mathcal{R}, \mathcal{S}}$. This motivates the following definition:

Definition 3. For any subset $\mathcal{S} \subseteq \mathcal{X}_{\text{in}}$, the set of connected components of $G_{\mathcal{R}, \mathcal{S}}$ is called an \mathcal{R} -robustness structure.

Let \mathbf{B} be an \mathcal{R} -robustness structure, and let $\mathcal{S} = \cup \mathbf{B}$. Let $f_{\mathbf{B}} : \mathcal{S} \rightarrow \mathbf{B}$ be the map that maps each $x \in \mathcal{S}$ to the corresponding block of \mathbf{B} containing x . Any stochastic map κ that is \mathcal{R} -robust on \mathcal{S} factorizes through $f_{\mathbf{B}}$, in the sense that there is a stochastic map κ' that maps each block in \mathbf{B} to a probability distribution on \mathcal{X}_0 and that satisfies $\kappa = \kappa' \circ f_{\mathbf{B}}$. Conversely, any stochastic map κ that factorizes through $f_{\mathbf{B}}$ is \mathcal{R} -robust.

To any joint probability distribution p_{in} on \mathcal{X}_{in} with $p(X_{\text{in}} \in \mathcal{S}) = 1$ we can associate a random variable $B = f_{\mathbf{B}}(X_1, \dots, X_n)$. If κ is \mathcal{R} -robust on \mathcal{S} , then X_0 is independent of X_1, \dots, X_n given B . Note that the random variable B is only defined on $\cup \mathbf{B}$, which is a set of measure one with respect to p_{in} . The situation is illustrated by the following graph:



When the robustness specification \mathcal{R} is fixed, how much freedom is left to choose a robust stochastic map κ ? More precisely, how many components can an \mathcal{R} -robustness structure \mathbf{B} have?

Lemma 4. Let \mathbf{B} be a robustness structure of the robustness specification \mathcal{R} . Let $R \subseteq [n]$, $\mathcal{S} = [n] \setminus R$ and $\mathcal{Y}_R = \{x_R \in \mathcal{X}_R : (R, x_R) \in \mathcal{R}\}$. Then

$$|\mathbf{B}| \leq |\mathcal{Y}_R| + |\mathcal{X}_R \setminus \mathcal{Y}_R| \cdot |\mathcal{X}_{\mathcal{S}}|.$$

Proof. The set \mathcal{S} is the disjoint union of the $|\mathcal{Y}_R|$ sets $C(R, x_R) \cap \mathcal{S}$ for $x_R \in \mathcal{Y}_R$ and the $|\mathcal{X}_R \setminus \mathcal{Y}_R| \cdot |\mathcal{X}_{\mathcal{S}}|$ singletons $\{x\} \subset \mathcal{S}$ with $x|_R \notin \mathcal{Y}_R$. Each of these sets induces a connected subgraph of $G_{\mathcal{R}}$. The statement now follows from Proposition 2. \square

Example 5. Suppose that $\mathcal{S} = \mathcal{X}_{\text{in}}$. This means that any \mathcal{R} -robustness structure \mathbf{B} satisfies $\cup \mathbf{B} = \mathcal{X}_{\text{in}}$. If $G_{\mathcal{R}}$ is connected, then \mathbf{B} has just a single block. In this case the bound of Lemma 4 is usually not tight. On the other hand, the bound is tight if $\mathcal{R} = \{(R, x_R) : x_R \in \mathcal{X}_R\}$.

Remark 6 (Relation to coding theory). Assume that all d_i are equal. We can interpret \mathcal{X}_{in} as the set of words of length n over the alphabet $[d_1]$. Consider the uniform case $\mathcal{R} = \mathcal{R}_k$. Then the task is to find a collection of subsets such that any two different subsets have Hamming distance at least $n - k + 1$. A related problem appears in coding theory: A code is a subset \mathcal{Y} of \mathcal{X}_{in} and corresponds to the case that each element of \mathbf{B} is a singleton. If distinct elements of the code have Hamming distance at least $n - k + 1$, then a message can be reliably decoded even if only k letters are transmitted. If all letters are transmitted, but up to k letters may contain an error, then this error may at least be detected; hence such codes are called *error detecting codes*. In this setting, the function $f_{\mathbf{B}}$ can be interpreted as the decoding operation. The problem of finding a largest possible code such that all code words have a fixed minimum distance is also known as the *sphere packing problem*. The maximal size $A_{d_1}(n, n - k + 1)$ of such a code is unknown in general.

3. CANALYZING FUNCTIONS

Our notion of \mathcal{R} -robust functional modalities naturally generalizes and is motivated by canalizing [9] and nested canalizing functions [10]. Let $f : \mathcal{X}_{\text{in}} \rightarrow \mathcal{X}_0$ be a function, also called (deterministic) map. Such a map can be considered as a special case of a stochastic map by identifying f with

$$\kappa^f(x; x_0) := \begin{cases} 1, & \text{if } f(x) = x_0 \\ 0, & \text{otherwise} \end{cases}.$$

We say that f is (R, x_R) -canalizing, if the value of f does not depend on the input variables $X_{[n] \setminus R}$ given that the input variables X_R are in state x_R . In other words, an (R, x_R) -canalizing function is assumed to be constant on the corresponding cylinder set:

$$x, x' \in C(R, x_R) \quad \Rightarrow \quad f(x) = f(x').$$

Given a robustness specification \mathcal{R} , we say that a function f is \mathcal{R} -canalizing if it is (R, x_R) -canalizing for all $(R, x_R) \in \mathcal{R}$. Clearly, the set of \mathcal{R} -canalizing functions strongly depends on \mathcal{R} . On one hand, any function is \mathcal{R} -canalizing with respect to

$$\mathcal{R} = \{([n], x) : x \in \mathcal{X}_{\text{in}}\}.$$

On the other hand, for two different elements $i, j \in [n]$, and

$$\mathcal{R} = (\{i\} \times \mathcal{X}_i) \cup (\{j\} \times \mathcal{X}_j),$$

any \mathcal{R} -canalizing function is constant. Note that constant functions are \mathcal{R} -canalizing for any \mathcal{R} .

The following statement directly follows from Proposition 2:

Proposition 7. *A function $f : \mathcal{X}_{\text{in}} \rightarrow \mathcal{X}_0$ is \mathcal{R} -canalizing if and only if κ^f is \mathcal{R} -robust in $\mathcal{S} = \mathcal{X}_{\text{in}}$.*

Particular cases of canalizing functions have been studied in the context of robustness:

Example 8. (1) *Canalizing functions.* A function f with domain \mathcal{X}_{in} is *canalizing* in the sense of [9], if there exist an input node $k \in [n]$, an input value $a \in \mathcal{X}_k$, and an output value $b \in \mathcal{X}_0$ such that the value of f is independent of $x_{[n] \setminus \{k\}}$, given that

$x|_k = a$. In other words, $f(x) = f(y) = b$ whenever $x|_k = y|_k = a$. A canalizing function is \mathcal{R} -canalizing with

$$\mathcal{R} := \{(R, x_R) : R \subseteq [n], k \in R, x_R \in \mathcal{X}_R, x_R|_k = a\}.$$

(2) *Nested canalizing functions* have been studied in [10]. A function f is *nested canalizing* in the variable order X_1, \dots, X_n with canalizing input values $a_1 \in \mathcal{X}_1, \dots, a_n \in \mathcal{X}_n$ and canalized output values b_1, \dots, b_n if f satisfies $f(x) = b_k$ for all $x \in \mathcal{X}$ satisfying $x|_k = a_k$ and $x|_i \neq a_i$ for all $i < k$. Let $\mathcal{R} := \biguplus_{k=1}^n \mathcal{R}^{(k)}$, where

$$\mathcal{R}^{(k)} := \{(R, x_R) : [k] \subseteq R, x_R|_1 \neq a_1, \dots, x_R|_{k-1} \neq a_{k-1}, x_R|_k = a_k\}.$$

It is easy to see that f is a nested canalizing function if and only if it is \mathcal{R} -canalizing.

The set of Boolean nested canalizing functions has been described algebraically in [7] as a variety over the finite field \mathbb{F}_2 . Here, we use a different viewpoint, which allows to study not only deterministic functions, but also stochastic functions.

4. ROBUSTNESS AND GIBBS REPRESENTATION

Let (κ_A) be a collection of functional modalities, as defined in Section 2. Instead of providing a list of all functional modes κ_A , one can describe them in more mechanistic terms. To illustrate this, we first consider an example from the field of neural networks: Assume that the output node receives an input $x = (x_1, \dots, x_n) \in \{-1, +1\}^n$ and generates the output $+1$ with probability

$$\kappa(x_1, \dots, x_n; +1) := \frac{1}{1 + e^{-(\sum_{i=1}^n w_i x_i - \eta)}}.$$

For an arbitrary output x_0 this implies

$$(3) \quad \kappa(x_1, \dots, x_n; x_0) := \frac{e^{\frac{1}{2}(\sum_{i=1}^n w_i x_i - \eta)x_0}}{e^{\frac{1}{2}(\sum_{i=1}^n w_i x_i - \eta)(-1)} + e^{\frac{1}{2}(\sum_{i=1}^n w_i x_i - \eta)(+1)}}.$$

The structure of this representation of the stochastic map κ already suggests what the function should be after a knockout of a set S of input nodes: Simply remove the contribution of all the nodes in S . The post-knockout function is then given by

$$(4) \quad \kappa_R(x_R; +1) := \frac{e^{\frac{1}{2}(\sum_{i \in R} w_i x_i - \eta)x_0}}{e^{+\frac{1}{2}(\sum_{i \in R} w_i x_i - \eta)} + e^{-\frac{1}{2}(\sum_{i \in R} w_i x_i - \eta)}},$$

where $R = [n] \setminus S$. These post-knockout functional modalities are based on the decomposition of the sum that appears in (3).

More generally, we consider the following model of (κ_A) :

$$(5) \quad \kappa_A(x_A; x_0) = \frac{e^{\sum_{B \subseteq A} \phi_B(x_A|_B, x_0)}}{\sum_{x_0'} e^{\sum_{B \subseteq A} \phi_B(x_A|_B, x_0')}},$$

where the ϕ_B are functions on $\mathcal{X}_B \times \mathcal{X}_0$. Such a sum decomposition of κ is referred to as a *Gibbs representation* of κ and contains more information than κ itself. Clearly, each κ_A is strictly positive. Using the Möbius inversion, it is easy to see that each strictly positive family (κ_A) has a representation of the form (5) with

$$(6) \quad \phi_A(x_A, x_0) := \sum_{C \subseteq A} (-1)^{|A \setminus C|} \ln \kappa_C(x_A|_C; x_0).$$

Note that this representation is not unique: If an arbitrary function of x_A (that does not depend on x_0) is added to the function ϕ_A , then the function κ_A , defined via (5), does not change.

A single robustness constraint has the following consequences for the ϕ_A .

Proposition 9. *Let $S \subseteq [n]$ and $R = [n] \setminus S$, and let (κ_A) be strictly positive functional modalities with Gibbs potentials (ϕ_A) . Then (κ_A) is robust in x against knockout of S if and only if $\sum_{B \subseteq [n], B \not\subseteq R} \phi_B(x|_B, x_0)$ does not depend on x_0 .*

Proof. Denote by $\tilde{\phi}_A$ the potentials defined via (6). Then (1) is equivalent to

$$\sum_{B \subseteq [n]} \tilde{\phi}_B(x|_B, x_0) = \sum_{B \subseteq R} \tilde{\phi}_B(x|_B, x_0) \iff \sum_{\substack{B \subseteq [n] \\ B \not\subseteq R}} \tilde{\phi}_B(x|_B, x_0) = 0.$$

The statement follows from the fact that $\phi_B(x|_B; x_0) - \tilde{\phi}_B(x|_B; x_0)$ is independent of x_0 (for fixed x). \square

Example 10. Consider $n = 2$ binary inputs, $\mathcal{X}_1 = \mathcal{X}_2 = \{0, 1\}$, and let $\mathcal{S} = \{(0, 0), (1, 1)\}$. Then 1-robustness on \mathcal{S} means

$$\kappa_{\{1\}}(x_1; x_0) = \kappa_{\{1,2\}}(x_1, x_2; x_0) = \kappa_{\{2\}}(x_2; x_0)$$

for all x_0 whenever $x_1 = x_2$. By Proposition 9 this translates into the conditions

$$(7) \quad \phi_{\{1,2\}}(x_1, x_2; x_0) + \phi_{\{1\}}(x_1; x_0) = 0 = \phi_{\{1,2\}}(x_1, x_2; x_0) + \phi_{\{2\}}(x_2; x_0)$$

for all x_0 whenever $x_1 = x_2$ for the potentials (ϕ_A) defined via (6). This means: Assuming that (κ_A) is 1-robust, it suffices to specify the four functions

$$\phi_\emptyset(x_0), \phi_{\{1\}}(x_1; x_0), \phi_{\{1,2\}}(0, 1; x_0), \phi_{\{1,2\}}(1, 0; x_0).$$

The remaining potentials can be deduced from (7). If only the values of (κ_A) for $x \in \mathcal{S}$ are needed, then it suffices to specify $\phi_\emptyset(x_0)$ and $\phi_{\{1\}}(x_1; x_0)$.

Does \mathcal{R} -robustness in x imply any structural constraints on (κ_A) ? If (κ_A) is \mathcal{R} -robust in x for all x belonging to a set \mathcal{S} , then the corresponding conditions imposed by Proposition 9 depend on \mathcal{S} . In this section, we are interested in conditions that are independent of \mathcal{S} . Such conditions allow to define sets of functional modalities that contain all \mathcal{R} -robust functional modalities for all possible sets \mathcal{S} . If \mathcal{S} (which will be the support of the input distribution in Section 5) is unknown from the beginning, then the system can choose its policy within such a restricted set of functional modalities. To find results that are independent of \mathcal{S} , our trick is to find a set $\tilde{M}_{\mathcal{R}}$ of functional modalities such that (κ_A) can be approximated on \mathcal{S} by functional modalities in $\tilde{M}_{\mathcal{R}}$. The approximation will be independent of \mathcal{S} .

We first consider the special case $\mathcal{R} = \mathcal{R}_k := \{(R, x_R) : R \subseteq [n], |R| \geq k, x_R \in \mathcal{X}_R\}$. For simplicity, we replace any prefix or subscript \mathcal{R}_k by k . Denote by M_{k+1} the set of all functional modalities (κ_A) such that there exist potentials ϕ_A of the form

$$\phi_A(x_A; x_0) = \sum_{\substack{B \subseteq A \\ |B| \leq k+1}} \alpha_{A,B} \Psi_B(x_A|_B; x_0),$$

where $\alpha_{A,B} \in \mathbb{R}$ and Ψ_B is an arbitrary function $\mathbb{R}^{\mathcal{X}_B \times \mathcal{X}_0} \rightarrow \mathbb{R}$. The set M_{k+1} is called the *family of $(k+1)$ -interaction functional modalities*. Note that the functions Ψ_B do not depend on A . This ensures a certain interdependence among the functional

modalities κ_A . The name “ $(k+1)$ -interaction” comes from the fact that each potential Ψ_B depends on the k (or less) variables in B plus the output variable X_0 . Since M_{k+1} only contains strictly positive functional modalities, we are also interested in the *closure* of M_{k+1} with respect to the usual topology on the space of matrices, considered as elements of a finite-dimensional real vector space.

Example 11. The functional modalities (4), derived from the classical model (3) of a neural network, belong to M_2 . To illustrate the difference between M_2 and its closure, consider the functional modalities (κ_A) with

$$\kappa_A(x_1, \dots, x_n; x_0) := \frac{e^{\frac{\beta}{2}(\sum_{i \in A} w_i x_i - \eta) x_0}}{e^{-\frac{\beta}{2}(\sum_{i \in A} w_i x_i - \eta)} + e^{+\frac{\beta}{2}(\sum_{i \in A} w_i x_i - \eta)}}.$$

If w_1, \dots, w_n and η are fixed and $\beta \rightarrow \infty$, then

$$(8) \quad \kappa_A(x_1, \dots, x_n; +1) \rightarrow \theta\left(\sum_{i \in A} w_i x_i - \eta\right),$$

where

$$\theta(x) = \begin{cases} +1, & \text{if } x > 0, \\ \frac{1}{2}, & \text{if } x = 0, \\ 0, & \text{if } x < 0. \end{cases}$$

The functional modalities (8) are deterministic limits of the probabilistic model (3), called *linear threshold functions*. They lie in the closure of M_2 , but not in M_2 itself.

Linear threshold functions are widely used as elementary building blocks in network dynamics, for example to build simple models of neural networks, metabolic networks or gene-regulation networks. Robustness against knockouts of such networks has been studied in [2], exploring the example of the yeast cell cycle.

Let \tilde{M}_{k+1} be the set of strictly positive functional modalities (κ_A) such that

$$(9) \quad \kappa_C(x_C; x_0) = \frac{1}{Z_{C, x_C}} \exp \left(\sum_{\substack{B \subseteq C \\ |B|=k}} \frac{1}{\binom{|C|}{k}} \ln(\kappa_B(x_C|_B; x_0)) \right) \\ = \frac{1}{Z_{C, x_C}} \left(\prod_{\substack{B \subseteq C \\ |B|=k}} \kappa_B(x_C|_B; x_0) \right)^{1/\binom{|C|}{k}}$$

for all $C \subseteq [n]$ with $|C| > k$, where Z_{C, x_C} is a normalization constant that ensures that $\kappa_C(x_C)$ is a probability distribution. Note that equations (9) can be used to parametrize the set \tilde{M}_{k+1} : The stochastic maps κ_A with $|A| \leq k$ can be chosen arbitrarily, while all other stochastic maps κ_C with $|C| > k$ can be computed by normalizing the geometric mean of the stochastic maps κ_B for $B \subseteq C$ and $|B| = k$.

Lemma 12. \tilde{M}_{k+1} is a subset of M_{k+1} . It consists of those functional modalities (κ_A) where the coefficients $\alpha_{A,B}$ additionally satisfy

$$(-1)^{|A|} \alpha_{A,B} = (-1)^{|A'|} \alpha_{A',B}, \quad \text{whenever } B \subseteq A \cap A' \text{ and } |B| < k,$$

and

$$(-1)^{|A'|} \alpha_{A',B} = (-1)^{|A|} \alpha_{A,B}, \quad \text{whenever } B \subseteq A \cap A' \text{ and } |B| = k.$$

for all $x_B \in \mathcal{X}_B$ and $x_0 \in \mathcal{X}_0$.

Proof. Assume that the coefficients $\alpha_{A,B}$ of $(\kappa_A) \in M_{k+1}$ satisfy the conditions stated in the lemma. We may multiply all functions Ψ_B by scalars and assume

$$(10) \quad \alpha_{A,B} = (-1)^{|A|-|B|}, \quad \text{if } |B| < k, \quad \alpha_{A,B} = (-1)^{|A|-k} \frac{k}{|A|}, \quad \text{if } |B| = k.$$

Then $\ln(\kappa_C(x_C; x_0))$ equals the logarithm of the normalization constant plus

$$\begin{aligned} & \sum_{A \subseteq C} \left(\sum_{\substack{B \subseteq A \\ |B| < k}} (-1)^{|A|-|B|} \Psi_B(x_C|_B; x_0) + \sum_{\substack{B \subseteq A \\ |B| = k}} (-1)^{|A|-k} \frac{k}{|A|} \Psi_B(x_C|_B; x_0) \right) \\ &= \sum_{\substack{B \subseteq C \\ |B| < k}} \left(\sum_{R \subseteq C \setminus B} (-1)^{|R|} \right) \Psi_B(x_C|_B; x_0) \\ & \quad + \sum_{\substack{B \subseteq C \\ |B| = k}} \left(\sum_{R \subseteq C \setminus B} (-1)^{|R|} \frac{k}{|R| + k} \right) \Psi_B(x_C|_B; x_0) \\ &= \sum_{\substack{B \subseteq C \\ |B| < k}} \left(\sum_{l=0}^{|C|-|B|} (-1)^l \binom{|C|-|B|}{l} \right) \Psi_B(x_C|_B; x_0) \\ & \quad + \sum_{\substack{B \subseteq C \\ |B| = k}} \left(\sum_{l=0}^{|C|-k} (-1)^l \binom{|C|-k}{l} \frac{k}{l+k} \right) \Psi_B(x_C|_B; x_0) \\ (11) \quad &= \sum_{\substack{B \subseteq C \\ |B| < k}} \delta_{|C|,|B|} \Psi_B(x_C|_B; x_0) + \sum_{\substack{B \subseteq C \\ |B| = k}} \frac{1}{\binom{|C|}{k}} \Psi_B(x_C|_B; x_0), \end{aligned}$$

where the identity $\sum_{i=0}^r \binom{r}{i} \frac{(-1)^i}{m+i} = 1 / \left((m+r) \binom{r+m-1}{m-1} \right)$ was used and $\delta_{a,b}$ denotes Kronecker's delta. For $|C| > k$ the first sum is empty, and it follows that κ_C satisfies the defining equality of \tilde{M}_{k+1} .

Conversely, if $(\kappa_A) \in \tilde{M}_{k+1}$, then let $\alpha_{A,B}$ be as in (10), and let

$$\Psi_B(x_B; x_0) = \log(\kappa_B(x_B; x_0)), \quad \text{for all } x_0 \in \mathcal{X}_0, x_B \in \mathcal{X}_B, |B| \leq k.$$

These coefficients $\alpha_{A,B}$ and functions Ψ_B together define an element $(\tilde{\kappa}_A) \in M_{k+1}$. The calculation (11) shows that

$$\tilde{\kappa}_A(x_A; x_0) = \begin{cases} \frac{1}{Z_{A,x_A}} \exp(\Psi_A(x_A; x_0) = \kappa_A(x_A; x_0)), & \text{if } |A| \leq k, \\ \frac{1}{Z_{A,x_A}} \exp\left(\sum_{\substack{B \subseteq A \\ |B| = k}} \frac{1}{\binom{|A|}{k}} \ln(\kappa_B(x_A|_B; x_0))\right), & \text{if } |A| > k, \end{cases}$$

and so $(\kappa_A) = (\tilde{\kappa}_A)$ belongs to M_{k+1} and is of the desired form. \square

Theorem 13. *Let (κ_A) be functional modalities. Then there exist functional modalities $(\tilde{\kappa}_A)$ in the closure of \tilde{M}_{k+1} such that the following holds: If (κ_A) is k -robust on a set $\mathcal{S} \subseteq \mathcal{X}_{\text{in}}$, then $\kappa_A(x|_A) = \tilde{\kappa}_A(x|_A)$ for all $A \subseteq [n]$ and all $x \in \mathcal{S}$. In particular, $(\tilde{\kappa}_A)$ belongs to the closure of the family of $(k+1)$ -interactions.*

Proof. Define $(\tilde{\kappa}_A)$ via

$$\tilde{\kappa}_A(x_A; x_0) = \begin{cases} \kappa_A(x_A; x_0), & \text{if } |A| \leq k, \\ \frac{1}{Z_{A, x_A}} \left(\prod_{\substack{B \subseteq A \\ |B|=k}} \kappa_B(x_A|_B; x_0) \right)^{1/\binom{|A|}{k}}, & \text{else,} \end{cases}$$

where Z_{A, x_A} is a normalization constant. By definition, $(\tilde{\kappa}_A)$ lies in the closure of \tilde{M}_{k+1} . Let $x \in \mathcal{S}$ and $C \subseteq [n]$. If $|C| \leq k$, then $\tilde{\kappa}_C(x|_C) = \kappa_C(x|_C)$ by definition of $\tilde{\kappa}_A$. So assume that $|C| > k$. By definition of k -robustness, if $x \in \mathcal{S}$, then $\kappa_C(x|_C) = \kappa_B(x|_B)$ for all $B \subset C$ with $|B| = k$. Therefore, if $x \in \mathcal{S}$ and $|C| > k$, then

$$\kappa_C(x|_C; x_0) = \left(\prod_{\substack{B \subseteq C \\ |B|=k}} \kappa_B(x_C|_B; x_0) \right)^{1/\binom{|C|}{k}}.$$

Therefore, if $x \in \mathcal{S}$ and $|C| > k$, then $Z_{C, x|_C} = 1$ and $\kappa_C(x|_C) = \tilde{\kappa}_C(x|_C)$. \square

Since M_{k+1} and \tilde{M}_{k+1} are independent of \mathcal{S} , Theorem 13 shows that these two families can be used to construct robust systems, when the set \mathcal{S} is not known a priori but must be learnt by the system, or when \mathcal{S} changes with time and the system must adapt.

If we are not interested in all functional modalities but just the stochastic map κ describing the unperturbed system, we can describe κ in terms of low interaction order. The family of $(k+1)$ -interaction stochastic maps, denoted by K_{k+1} , consists of all strictly positive maps κ such that

$$\ln \kappa(x; x_0) = \sum_{\substack{A \subseteq [n] \\ |A| \leq k}} \Psi_A(x|_A; x_0)$$

for some real functions $\Psi_A : \mathcal{X}_A \rightarrow \mathbb{R}$.

Corollary 14. *Let κ be a stochastic map. For given k there exists a stochastic map $\tilde{\kappa}$ in the closure of K_{k+1} such that the following holds: If κ is k -robust on a set \mathcal{S} , then $\kappa(x) = \tilde{\kappa}(x)$ for all $x \in \mathcal{S}$.*

Proof. If κ is k -robust on \mathcal{S} , there exist functional modalities $(\kappa_A)_A$ with $\kappa = \kappa_{[n]}$. Choose $(\tilde{\kappa}_A)$ as in Theorem 13. If $x \in \mathcal{S}$, then $\kappa(x) = \kappa_{[n]}(x) = \tilde{\kappa}_{[n]}(x)$. Hence the Corollary holds true with $\tilde{\kappa} = \tilde{\kappa}_{[n]}$. \square

Example 15. The functional modalities (4) do not lie in \tilde{M}_2 . This does not mean that neural networks are not robust: In fact, it is possible to naturally redefine the functional modalities (4) such that the new functional modalities lie in \tilde{M}_2 .

The construction (4) identifies the summand $w_i x_i x_0$ with $\phi_{\{i\}}$. Now we will make another identification: For each $i \in [n]$ let

$$\kappa_{\{i\}}(x_i; x_0) = \frac{1}{Z_{i, x_i}} \exp(n w_i x_i x_0 - \eta).$$

The unique extension of these stochastic maps to functional modalities (κ_A) in M_2 is given by

$$(12) \quad \kappa_A(x|_A; x_0) = \frac{1}{Z'_{A, x|_A}} \left(\prod_{i \in A} \kappa_{\{i\}}(x_i; x_0) \right)^{1/|A|} = \frac{1}{Z_{A, x|_A}} \exp \left(\frac{n}{|A|} \sum_{i \in A} w_i x_i x_0 - \eta \right),$$

where $Z_{A, x|_A}$ and $Z'_{A, x|_A}$ are constants determined by normalization. The functional modalities defined in this way lie in \tilde{M}_2 , and the stochastic map $\kappa_{[n]}$ agrees with (3). Note that, by tuning the parameters w_1, \dots, w_n , any combination of stochastic maps is possible for $\kappa_1, \dots, \kappa_n$. This shows that any element of \tilde{M}_2 has a representation of the form (12).

As in example 11 we can scale the weights w_i and the threshold η by a factor of β and send $\beta \rightarrow +\infty$. This leads to the rule

$$(13) \quad \kappa_A(x_A; +1) \rightarrow \theta\left(\frac{n}{|A|} \sum_{i \in A} w_i x_i - \eta\right),$$

which is a normalized variant of (8).

The rule (12) implements a renormalization of the effect of the remaining inputs under knockout. Similar renormalization procedures are sometimes used when training neural networks using Hebb's rule. Usually the total sum of the weights $\sum_i w_i$ is normalized to not grow to infinity. The rule (12) suggests that under knockout all remaining weights are amplified by a common factor.

The ideas leading to Theorem 13 can be applied to more general robustness structures \mathcal{R} as follows: For any $x \in X$ let

$$\mathcal{R}_x := \begin{cases} \{R \subseteq [n] : (R, x|_R) \in \mathcal{R}\}, & \text{if there exists } R \subseteq [n] \text{ with } (R, x|_R) \in \mathcal{R}, \\ \{[n]\}, & \text{else,} \end{cases}$$

and let \mathcal{R}_x^{\min} be the subset of inclusion-minimal elements of \mathcal{R}_x . If (κ_A) is \mathcal{R} -robust in \mathcal{S} , then

$$\kappa(x; x_0) = \kappa_R(x|_R; x_0) \quad \text{for any } R \in \mathcal{R}_x^{\min}, \quad x \in \mathcal{S}$$

and hence

$$\kappa(x; x_0) = \left(\prod_{R \in \mathcal{R}_x^{\min}} \kappa_R(x|_R; x_0) \right)^{1/|\mathcal{R}_x^{\min}|}.$$

For any $C \subseteq [n]$ let $\mathcal{R}_x^{\min}(C) = \{R \in \mathcal{R}_x^{\min} : C \subseteq R\}$. If \mathcal{R} is coherent, then we can deduce

$$(14) \quad \kappa_C(x|_C; x_0) = \left(\prod_{R \in \mathcal{R}_x^{\min}(C)} \kappa_R(x|_R; x_0) \right)^{1/|\mathcal{R}_x^{\min}(C)|}$$

for all $x \in \mathcal{S}$ with $\mathcal{R}_x^{\min}(C) \neq \emptyset$. This motivates the following definition: Denote by $\tilde{M}_{\mathcal{R}}$ the set of all strictly positive functional modalities that satisfy

$$\kappa_C(x|_C; x_0) = \frac{1}{Z_{C, x|_C}} \left(\prod_{R \in \mathcal{R}_x^{\min}(C)} \kappa_R(x|_R; x_0) \right)^{1/|\mathcal{R}_x^{\min}(C)|}$$

for all $x \in X$ and all $C \subseteq [n]$ with $\mathcal{R}_x^{\min}(C) \neq \emptyset$, where $Z_{C, x|_C}$ is a suitable normalization constant. The same proof as for Theorem 13 implies:

Theorem 16. *Let (κ_A) be functional modalities, and assume that \mathcal{R} is coherent. Then there exist functional modalities $(\tilde{\kappa}_A)$ in the closure of $\tilde{M}_{\mathcal{R}}$ such that the following holds: If (κ_A) is \mathcal{R} -robust on a set $\mathcal{S} \subseteq X$, then*

$$\kappa_A(x|_A) = \tilde{\kappa}_A(x|_A) \quad \text{for all } x \in \mathcal{S}.$$

As a generalization of Lemma 12, we can also describe $\tilde{\mathcal{M}}_{\mathcal{R}}$ as a set of functional modalities with limited interaction order. To simplify the presentation, we assume that \mathcal{R} is *saturated*, by which we mean the following: If $(R, x_R) \in \mathcal{R}$ for some $x_R \in \mathcal{X}_R$, then $(R, x'_R) \in \mathcal{R}$ for all $x_R \in \mathcal{X}_R$. In other words, a saturated robustness specification is given by enumerating a family of subsets of $[n]$. For example, the robustness structures \mathcal{R}_k are saturated, while the robustness structures defining canalizing and nested canalizing functions (see Section 3) are not saturated. If \mathcal{R} is saturated, then \mathcal{R}_x and \mathcal{R}_x^{\min} are independent of $x \in \mathcal{X}$.

Consider the family

$$\Delta = \{C \subseteq [n] : C \subseteq R \text{ for some } R \in \mathcal{R}_x^{\min} \text{ and } x \in \mathcal{X}\},$$

and let $\Delta(C) = \{R \in \Delta : R \subseteq C\}$. Let M_{Δ} be the set of all functional modalities (κ_A) such that there exist potentials ϕ_A of the form

$$(15) \quad \phi_A(x_A; x_0) = \sum_{B \in \Delta(A)} \alpha_{A,B} \Psi_B(x_A|_B; x_0),$$

where $\alpha_{A,B} \in \mathbb{R}$ and Ψ_B is an arbitrary function $\mathbb{R}^{\mathcal{X}_B \times \mathcal{X}_0} \rightarrow \mathbb{R}$. We call M_{Δ} the *family of Δ -interaction functional modalities*. Note that the functions Ψ_B do not depend on A . This ensures a certain interdependence among the functional modalities κ_A .

Lemma 17. *Assume that \mathcal{R} is coherent and saturated. $\tilde{\mathcal{M}}_{\mathcal{R}}$ is a subset of M_{Δ} .*

Proof. If $\mathcal{R}_x = \emptyset$, then Δ contains all sets. The Möbius inversion formula shows that M_{Δ} contains all strictly positive functional modalities. Therefore, we may assume that $\mathcal{R}_x \neq \emptyset$.

Define Gibbs potentials using the Möbius inversion (6). If $x \in \mathcal{S}$ and A is large enough such that $\mathcal{R}_x^{\min}(A) \neq \emptyset$, then

$$\begin{aligned} \sum_{\substack{C \subseteq A \\ C \in \mathcal{R}_x}} (-1)^{|A \setminus C|} \ln \kappa_C(x|_C; x_0) &= \sum_{\substack{C \subseteq A \\ C \in \mathcal{R}_x}} (-1)^{|A \setminus C|} \frac{1}{|\mathcal{R}_x^{\min}(C)|} \sum_{B \in \mathcal{R}_x^{\min}(C)} \ln \kappa_C(x|_C; x_0) \\ &= \sum_{\substack{C \subseteq A \\ C \in \mathcal{R}_x}} (-1)^{|A \setminus C|} \frac{1}{|\mathcal{R}_x^{\min}(C)|} \sum_{B \in \mathcal{R}_x^{\min}(C)} \ln \kappa_B(x|_B; x_0) \\ &= \sum_{B \in \mathcal{R}_x^{\min}(C)} \left\{ \sum_{R \subseteq A \setminus B} (-1)^{|A| - |R| - k} \frac{1}{|\mathcal{R}_x^{\min}(B \cup R)|} \right\} \ln \kappa_B(x|_B; x_0). \end{aligned}$$

Together with (6) this gives

$$\phi_A(x|_A, x_0) = \sum_{\substack{C \subseteq A \\ C \notin \mathcal{R}_x}} \alpha_{A,C} \ln \kappa_C(x|_C; x_0) + \sum_{\substack{C \subseteq A \\ C \in \mathcal{R}_x^{\min}}} \alpha_{A,C} \ln \kappa_C(x|_C; x_0),$$

where

$$\alpha_{A,C} = \begin{cases} (-1)^{|A| - |C|}, & \text{if } C \notin \mathcal{R}_x, \\ \sum_{R \subseteq A \setminus B} (-1)^{|A| - |R| - k} \frac{1}{|\mathcal{R}_x^{\min}(B \cup R)|} & \text{if } C \in \mathcal{R}_x^{\min}. \end{cases}$$

This is clearly of the form (15). \square

In the case $\mathcal{R} = \mathcal{R}_k$ the sum $\sum_{R \subseteq A \setminus B} (-1)^{|A| - |R| - k} \frac{1}{|\mathcal{R}_x^{\min}(B \cup R)|}$ that appears in the proof of Lemma 17 can be solved explicitly, resulting in the statement of Lemma 12. In the general case this is not possible.

Corollary 14 also generalizes. Let Δ be as above. The set K_Δ of Δ -interactions stochastic maps consists of all strictly positive stochastic maps κ such that

$$\ln \kappa(x; x_0) = \sum_{A \in \Delta} \Psi_A(x|_A; x_0)$$

for some real functions $\Psi_A : \mathcal{X}_A \rightarrow \mathbb{R}$.

Corollary 18. *Let κ be a stochastic map, and let \mathcal{R} be a coherent and saturated robustness specification. There exists a stochastic map $\tilde{\kappa}$ in the closure of K_Δ such that the following holds: If κ is \mathcal{R} -robust on a set \mathcal{S} , then $\kappa(x) = \tilde{\kappa}(x)$ for all $x \in \mathcal{S}$.*

The proof is the same as the proof of Corollary 14.

Remark 19. Instead of representing functional modalities as a family (κ_A) of stochastic maps, it is possible to use a single stochastic map $\hat{\kappa}$, operating on a larger space, that integrates the information from the family (κ_A) . The stochastic map $\hat{\kappa}$ can be constructed as follows: For each $i = 1, \dots, n$ let $\hat{\mathcal{X}}_i$ be the disjoint union of \mathcal{X}_i and one additional element, denoted by 0. This additional state represents the knockout of X_i . Let $\hat{\mathcal{X}}_{\text{in}} = \hat{\mathcal{X}}_1 \times \dots \times \hat{\mathcal{X}}_n$. For each $y \in \hat{\mathcal{X}}_{\text{in}}$ let $\text{supp}(y) = \{i : y_i \neq 0\}$. We define the stochastic map $\hat{\kappa} : \mathcal{X}_0 \times \hat{\mathcal{X}}_{\text{in}} \rightarrow \mathcal{X}$ via

$$\hat{\kappa}(x; x_0) = \kappa_{\text{supp}(x)}(x|_{\text{supp}(x)}; x_0).$$

This construction gives a one-to-one correspondence between functional modalities and stochastic maps from $\hat{\mathcal{X}}_{\text{in}}$ to \mathcal{X}_0 .

As an example, consider the functional modalities defined in (4). In this example, the construction of $\hat{\kappa}$ is particularly easy: It just amounts to extending the input space to $\{-1, 0, +1\}^n$. Equation (3) remains valid for $\hat{\kappa}$. The construction is more complicated for the functional modalities (12).

More generally, any Gibbs representation for functional modalities (κ_A) as in (5) extends to a Gibbs representation of $\hat{\kappa}$: For any $B \subseteq [n]$, $x_0 \in \mathcal{X}_0$ and $x \in \hat{\mathcal{X}}_{\text{in}}$ let

$$\hat{\phi}_B(x, x_0) = \begin{cases} \phi_B(x|_B, x_0), & \text{if } \text{supp}(x) \subseteq B, \\ 0, & \text{else.} \end{cases}$$

Then

$$\hat{\kappa}(x; x_0) = \frac{e^{\sum_{B \subseteq [n]} \hat{\phi}_B(x, x_0)}}{\sum_{x'_0 \in \mathcal{X}_0} e^{\sum_{B \subseteq [n]} \hat{\phi}_B(x, x'_0)}}.$$

5. ROBUSTNESS AND CONDITIONAL INDEPENDENCE

Given the probability distribution p_{in} of the input variables and a stochastic map κ describing the system, the joint probability distribution of the complete system can be computed from

$$p(x_0, x) = \kappa(x; x_0)p_{\text{in}}(x), \quad \text{for all } (x_0, x) \in \mathcal{X},$$

As shown in Proposition 2, robustness of stochastic maps is related to conditional independence constraints on the joint distribution. In this section we study the set of all joint distributions that arise from robust systems in this way.

Let \mathcal{R} be a robustness specification. By Proposition 2, the stochastic map κ is \mathcal{R} -robust on $\text{supp}(p_{\text{in}})$ if and only if for all $(R, x_R) \in \mathcal{R}$ the output X_0 is (stochastically) independent of $X_{[n] \setminus R}$, given that $X_R = x_R$. In the following, this conditional

independence (CI) statement will be written as $X_0 \perp\!\!\!\perp X_{[n]\setminus R} \mid X_R = x_R$. This motivates the following definition: A joint distribution p is called \mathcal{R} -robust if it satisfies $X_0 \perp\!\!\!\perp X_{[n]\setminus R} \mid X_R = x_R$ for all $(R, x_R) \in \mathcal{R}$. We denote by $\mathcal{P}_{\mathcal{R}}$ the set of all \mathcal{R} -robust probability distributions.

The single conditional independence statement $X_0 \perp\!\!\!\perp X_{[n]\setminus R} \mid X_R = x_R$ means that the conditional distributions satisfy

$$p(X_0 = x_0 \mid X_{\text{in}} = x) = p(X_0 = x_0 \mid X_R = x_R), \quad \text{for all } x \in X_{\text{in}} \text{ with } p(x) > 0 \\ \text{and } x|_R = x_R.$$

It is often convenient to use another definition that avoids the need to work with conditional distributions: The statement $X_0 \perp\!\!\!\perp X_{[n]\setminus R} \mid X_R = x_R$ holds if and only if

$$(16) \quad p(x_0, x_S, x_R)p(x'_0, x'_S, x_R) = p(x_0, x'_S, x_R)p(x'_0, x_S, x_R),$$

for all $x_0, x'_0 \in X_0, x_S, x'_S \in X_S$ and $x_R \in X_R$. Here, $p(x_0, x_S, x_R)$ is an abbreviation of $p(X_0 = x_0, X_S = x_S, X_R = x_R)$. It is not difficult to see that these two definitions of conditional independence are equivalent. The formulation in terms of determinantal equations is used in algebraic statistics [4] and will also turn out to be useful here.

A joint probability distribution p can be written as a $d_0 \times |X_{\text{in}}|$ -matrix. Each equation (16) imposes conditions on this matrix saying that certain submatrices have rank one. To be precise, for any edge (x, x') in the graph $G_{\mathcal{R}}$ (defined in Section 2) equations (16) for all $x_0, x'_0 \in X_0$ require that the submatrix $(p_{kz})_{k \in X_0, z \in \{x, x'\}}$ has rank one. For any $x \in X_{\text{in}}$ denote by \tilde{p}_x the vector with components $\tilde{p}_x(x_0) = p(X_0 = x_0, X_{\text{in}} = x)$ for $x_0 \in X_0$. Then a distribution p lies in $\mathcal{P}_{\mathcal{R}}$ if and only if \tilde{p}_x and \tilde{p}_y are proportional for all edges (x, y) of $G_{\mathcal{R}}$. Observe that \tilde{p}_x and \tilde{p}_y are proportional if and only if either (i) one of \tilde{p}_x and \tilde{p}_y vanishes or (ii) $\kappa(x) = \kappa(y)$. This observation allows to reformulate the equivalence (1) \Leftrightarrow (3) of Proposition 2 as follows:

Lemma 20. *Let $\mathcal{S} = \{x \in X_{\text{in}} : \tilde{p}_x \neq 0\}$. A distribution p lies in $\mathcal{P}_{\mathcal{R}}$ if and only if \tilde{p}_x and \tilde{p}_y are proportional whenever $x, y \in \mathcal{S}$ lie in the same connected component of $G_{\mathcal{R}, \mathcal{S}}$.*

For any family \mathbf{B} of subsets of X_{in} let $\mathcal{P}_{\mathbf{B}}$ be the set of probability distributions p on \mathcal{X} that satisfy the following two conditions:

- (1) $\cup \mathbf{B} = \{x \in X_{\text{in}} : \tilde{p}_x \neq 0\}$,
- (2) \tilde{p}_x and \tilde{p}_y are proportional, whenever there exists $\mathcal{Z} \in \mathbf{B}$ such that $x, y \in \mathcal{Z}$.

Then $\mathcal{P}_{\mathcal{R}} = \cup_{\mathbf{B}} \mathcal{P}_{\mathbf{B}}$, where the union is over all \mathcal{R} -robustness structures \mathbf{B} . The disadvantage of this decomposition is that there are \mathcal{R} -robustness structures \mathbf{B}, \mathbf{B}' such that $\mathcal{P}_{\mathbf{B}}$ is a subset of the topological closure $\overline{\mathcal{P}_{\mathbf{B}'}}$ of $\mathcal{P}_{\mathbf{B}'}$. In other words, each $p \in \mathcal{P}_{\mathbf{B}}$ can be approximated arbitrarily well by elements of $\mathcal{P}_{\mathbf{B}'}$, and therefore in many cases it suffices to only consider $\mathcal{P}_{\mathbf{B}'}$. The following definition is needed:

Definition 21. An \mathcal{R} -robustness structure \mathbf{B} is *maximal* if and only if $\cup \mathbf{B} := \cup_{\mathcal{Z} \in \mathbf{B}} \mathcal{Z}$ satisfies any of the following equivalent conditions:

- (1) For any $x \in X_{\text{in}} \setminus \cup \mathbf{B}$ there are edges $(x, y), (x, z)$ in $G_{\mathcal{R}}$ such that $y, z \in \cup \mathbf{B}$ do not lie in the same connected component of $G_{\mathcal{R}, \cup \mathbf{B}}$.
- (2) For any $x \in X_{\text{in}} \setminus \cup \mathbf{B}$ the induced subgraph $G_{\mathcal{R}, \cup \mathbf{B} \cup \{x\}}$ has fewer connected components than $G_{\mathcal{R}, \cup \mathbf{B}}$.

Lemma 22. $\mathcal{P}_{\mathcal{R}}$ equals the disjoint union $\bigcup_{\mathbf{B}} \mathcal{P}_{\mathbf{B}}$, where the union is over all \mathcal{R} -robustness structures. Alternatively, $\mathcal{P}_{\mathcal{R}}$ equals the (non-disjoint) union $\bigcup_{\mathbf{B}} \overline{\mathcal{P}_{\mathbf{B}}}$, where the union is over all maximal \mathcal{R} -robustness structures.

Proof. The first statement follows directly from the above considerations. To see that it suffices to take maximal \mathcal{R} -robustness structures in the second decomposition, consider an \mathcal{R} -robustness structure \mathbf{B} that is not maximal. By definition there exists $x \in \mathcal{X}_{\text{in}} \setminus \bigcup \mathbf{B}$ such that the induced subgraph $G_{\mathcal{R}, \bigcup \mathbf{B} \cup \{x\}}$ has at least as many connected components as $G_{\mathcal{R}, \bigcup \mathbf{B}}$. Let \mathbf{B}' be the family of connected components of $G_{\mathcal{R}, \bigcup \mathbf{B} \cup \{x\}}$. If $G_{\mathcal{R}, \bigcup \mathbf{B} \cup \{x\}}$ has the same number of connected components as $G_{\mathcal{R}, \bigcup \mathbf{B}}$, then there is $\mathcal{Y} \in \mathbf{B}$ such that $\mathcal{Y} \cup \{x\} \in \mathbf{B}'$, otherwise let $\mathcal{Y} \in \mathbf{B}$ be arbitrary. Let $y \in \mathcal{Y}$. For any $p \in \mathcal{P}_{\mathbf{B}}$ and $\epsilon > 0$ define a probability distribution p_{ϵ} via

$$p_{\epsilon}(x_0, z) = \begin{cases} p(x_0, z), & \text{if } z \notin \{x, y\}, \\ (1 - \epsilon)p(x_0, x), & \text{if } z = y, \\ \epsilon p(x_0, x), & \text{if } z = x. \end{cases}$$

Then $p_{\epsilon} \in \mathcal{P}_{\mathbf{B}'}$, and hence $\mathcal{P}_{\mathbf{B}} \subseteq \overline{\mathcal{P}_{\mathbf{B}'}}$. If \mathbf{B}' is not maximal, we may iterate the process. \square

The following lemma sheds light on the structure of $\overline{\mathcal{P}_{\mathbf{B}}}$:

Lemma 23. Fix an \mathcal{R} -robustness structure \mathbf{B} . Then $\overline{\mathcal{P}_{\mathbf{B}}}$ consists of all probability measures of the form

$$(17) \quad p(X_0 = x_0, X_{\text{in}} = x) = \begin{cases} \mu(\mathcal{Z}) \lambda_{\mathcal{Z}}(x) p_{\mathcal{Z}}(x_0), & \text{if } x \in \mathcal{Z} \in \mathbf{B}, \\ 0, & \text{if } x \in \mathcal{X}_{\text{in}} \setminus \bigcup \mathbf{B}, \end{cases}$$

where μ is a probability distribution on \mathbf{B} and $\lambda_{\mathcal{Z}}$ is a probability distribution on \mathcal{Z} for each $\mathcal{Z} \in \mathbf{B}$ and $(p_{\mathcal{Z}})_{\mathcal{Z} \in \mathbf{B}}$ is a family of probability distributions on \mathcal{X}_0 .

Proof. It is easy to see that (17) defines indeed a probability distribution. By Lemma 20 it belongs to $\mathcal{P}_{\mathbf{B}}$. In the other direction, any probability measure can be written as a product

$$p(x_0, x_1, \dots, x_n) = p(\mathcal{Z}) p(x_1, \dots, x_n | (X_1, \dots, X_n) \in \mathcal{Z}) p(x_0 | x_1, \dots, x_n),$$

if $(x_1, \dots, x_n) \in \mathcal{Z} \in \mathbf{B}$, and if p is an \mathcal{R} -robust probability distribution, then $p_{\mathcal{Z}}(x_0) := p(x_0 | x_1, \dots, x_n)$ depends only on the block \mathcal{Z} in which (x_1, \dots, x_n) lies. \square

Lemma 22 decomposes the set $\mathcal{P}_{\mathcal{R}}$ of robust probability distributions into the closures of the smooth manifolds $\mathcal{P}_{\mathbf{B}}$, where \mathbf{B} runs over the maximal \mathcal{R} -robustness structures. Lemma 23 gives natural parametrizations of these manifolds.

By comparison, Theorem 16 and Lemma 17 describe robustness from a different point of view. The result can be translated to the setting of this section as follows:

Corollary 24. Suppose that \mathcal{R} is a coherent and saturated robustness structure, and define Δ as in Section 4. If $p \in \mathcal{P}_{\mathbf{B}}$, then there exists a stochastic map $\tilde{\kappa} \in K_{\Delta}$ such that $p(x_0 | x) = \tilde{\kappa}(x; x_0)$ for all $x \in \bigcup \mathbf{B}$.

In the statement of the corollary note that $p(X_{\text{in}} = x) > 0$ for all $x \in \bigcup \mathbf{B}$, and hence the conditional distribution $p(x_0 | x)$ is well-defined in this case.

Corollary 24 can also be viewed from the perspective of hierarchical models: Let $\tilde{\Delta} = \{\{1, \dots, n\} \cup \{S \cup \{0\} : S \in \Delta\}\}$. The *hierarchical loglinear model* $\mathcal{E}_{\tilde{\Delta}}$ consists of all probability distributions p on \mathcal{X} of the form

$$\log(p(x)) = \sum_{A \subseteq \tilde{\Delta}} \tilde{\phi}_A(x|_A),$$

where ϕ_A is a real function with domain \mathcal{X}_A . By the results of this section, $\mathcal{E}_{\tilde{\Delta}}$ is a smooth manifold containing $\mathcal{P}_{\mathcal{R}}$ in its closure. See [11, 4] for more on hierarchical loglinear models.

Remark 25. It is also possible to derive the decomposition in Lemma 22 from results from commutative algebra. Since the equations (16) that describe conditional independence are algebraic, they generate a polynomial ideal, called *conditional independence ideal*. In this case the ideal is a *generalized binomial edge ideal*, as defined in [13]. For such ideals, the primary decomposition is known and corresponds precisely to the decomposition of the set of robust distributions as presented in Lemma 22. The parametrization of Lemma 23 can be considered as a surjective polynomial map and shows that all components of the decomposition are rational.

6. k -ROBUSTNESS

In this section we consider the symmetric case $\mathcal{R} = \mathcal{R}_k$. As above, we replace any prefix or subscript \mathcal{R} by k .

If $k = 0$, then any pair (x, y) is an edge in G_0 . This means that any 0-robustness structure \mathbf{B} contains only one set. There is only one maximal 0-robustness structure, namely $\overline{\mathbf{B}} = \{\mathcal{X}_{\text{in}}\}$. The set \mathcal{R}_0 is irreducible. This corresponds to the fact that \mathcal{P}_0 is defined by $X_0 \perp\!\!\!\perp X_{\text{in}}$.

$\overline{\mathbf{B}}$ is actually a maximal k -robustness structure for any $0 \leq k < n$. This illustrates the fact that the single CI statement $X_0 \perp\!\!\!\perp X_{\text{in}}$ implies all other CI statements of the form $X_0 \perp\!\!\!\perp X_{[n] \setminus R} \mid X_R = x_R$. The corresponding set $\mathcal{P}_{\overline{\mathbf{B}}}$ contains all probability distributions of \mathcal{P}_k of full support.

Now let $k = 1$. In the case $n = 2$ we obtain results by Alexander Fink, which can be reformulated as follows [5]: *Let $n = 2$. A 1-robustness structure \mathbf{B} is maximal if and only if the following statements hold:*

- Each $B \in \mathbf{B}$ is of the form $B = S_1 \times S_2$, where $S_1 \subseteq \mathcal{X}_1, S_2 \subseteq \mathcal{X}_2$.
- For every $x_1 \in \mathcal{X}_1$ there exists $B \in \mathbf{B}$ and $x_2 \in \mathcal{X}_2$ such that $(x_1, x_2) \in B$, and conversely.

In [5] a different description is given: The block $S_1 \times S_2$ can be identified with the complete bipartite graph on S_1 and S_2 . In this way, every maximal 1-robustness structure corresponds to a collection of complete bipartite subgraphs with vertices in $\mathcal{X}_1 \cup \mathcal{X}_2$ such that every vertex in \mathcal{X}_1 and \mathcal{X}_2 , respectively, is part of one such subgraph. Figure 2 shows an example.

This result generalizes in the following way:

Lemma 26. *A 1-robustness structure \mathbf{B} is maximal if and only if the following statements hold:*

- Each $B \in \mathbf{B}$ is of the form $B = S_1 \times \dots \times S_n$, where $S_i \subseteq \mathcal{X}_i$.
- $\bigcup_{S_1 \times \dots \times S_n \in \mathbf{B}} S_i = \mathcal{X}_i$ for all $i \in [n]$




FIGURE 2. A 1-robustness structures for two variables. a) The graph $G_{1,S}$. b) The representation in terms of bipartite graphs.

Proof. Suppose that \mathbf{B} is maximal. Let $\mathcal{Y} \in \mathbf{B}$ and let S_i be the projection of $\mathcal{Y} \subset \mathcal{X}_{\text{in}}$ to \mathcal{X}_i . Let $\mathcal{Y}' = S_1 \times \dots \times S_n$. Then $\mathcal{Y} \subseteq \mathcal{Y}'$. We claim that $(\mathbf{B} \setminus \{\mathcal{Y}\}) \cup \{\mathcal{Y}'\}$ is another 1-robustness structure with the same number of components as \mathbf{B} , and by maximality we can conclude $\mathcal{Y} = \mathcal{Y}'$. By Definition 3 we need to show that $G_{\mathcal{R}, \mathcal{Y}'}$ is connected and that $G_{\mathcal{R}, \mathcal{Z} \cup \mathcal{Y}'}$ is not connected for all $\mathcal{Z} \in \mathbf{B} \setminus \{\mathcal{Y}\}$. The first condition follows from the fact that $G_{\mathcal{R}, \mathcal{Y}}$ is connected. For the second condition assume to the contrary that there are $x \in \mathcal{Y}'$ and $y \in \mathcal{Z}$ such that $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ disagree in at most $n - 1$ components. Then there exists a common component $x_l = y_l$. By construction there exists $z = (z_1, \dots, z_n) \in \mathcal{Y}$ such that $z_l = y_l = x_l$, hence $\mathcal{Y} \cup \mathcal{Z}$ is connected, in contradiction to the assumptions. This shows that each \mathcal{Y} has a product structure.

Write $\mathcal{Y} = S_1^{\mathcal{Y}} \times \dots \times S_n^{\mathcal{Y}}$ for each $\mathcal{Y} \in \mathbf{B}$. Obviously $S_i^{\mathcal{Y}} \cap S_i^{\mathcal{Z}} = \emptyset$ for all $i \in [n]$ and all $\mathcal{Y}, \mathcal{Z} \in \mathbf{B}$ if $\mathcal{Y} \neq \mathcal{Z}$. For the second assertion, assume to the contrary that $l \in \mathcal{X}_i$ is contained in no $S_i^{\mathcal{Y}}$. Take any $\mathcal{Y} \in \mathbf{B}$ and define $\mathcal{Y}' := S_1^{\mathcal{Y}} \times \dots \times (S_i^{\mathcal{Y}} \cup \{l\}) \times \dots \times S_n^{\mathcal{Y}}$. Then $(\mathbf{B} \setminus \{\mathcal{Y}\}) \cup \{\mathcal{Y}'\}$ is another 1-robustness structure with the same number of components as \mathbf{B} , contradicting the assumptions.

Conversely, assume that \mathbf{B} is a 1-robustness structure satisfying the two assertions of the theorem. For any $x \in \mathcal{X}_{\text{in}} \setminus \cup \mathbf{B}$ there exist $y_1, \dots, y_n \in \cup \mathbf{B}$ such that $x_1 = y_1, \dots, x_n = y_n$. Since $x \notin \cup \mathbf{B}$ the points y_1, \dots, y_n cannot all belong to the same block of \mathbf{B} . If y_i and y_j belong to different blocks of \mathbf{B} , then the two edges (x, y_i) and (x, y_j) of G_1 show that \mathbf{B} is maximal. \square

The last result can be reformulated in terms of n -partite graphs generalizing [5]: Namely, the 1-robustness structures are in one-to-one relation with the n -partite subgraphs of M_{d_1, \dots, d_n} such that every connected component is itself a complete n -partite subgraph M_{e_1, \dots, e_n} with $e_i > 0$ for all $i \in [n]$. Here, an n -partite graph is a graph which can be coloured by n colours such that no two vertices with the same colour are adjacent.

Unfortunately the nice product form of the maximal 1-robustness structures does not generalize to $k > 1$:

Example 27 (Binary inputs). If $n = 3$ and $d_1 = d_2 = d_3 = 2$, then the graph G_2 is the graph of the cube. For a maximal 2-robustness structure \mathbf{B} the set $\mathcal{X}_{\text{in}} \setminus \cup \mathbf{B}$ can be any one of the following (see Fig. 3):

- The empty set
- A set of cardinality 4 corresponding to a plane leaving two connected components of size 2
- A set of cardinality 4 containing all vertices with the same parity.




FIGURE 3. The four symmetry classes of maximal 2-robustness structures of three binary inputs, see Example 27.




FIGURE 4. A maximal 3-robustness structure for four binary inputs.




FIGURE 5. The 2-robustness structure from Example 28. The graph G_2 is the graph of a hypercube of dimension four, where diagonals have been added to the two-dimensional faces. Only the edges of G_2 that connect vertices of Hamming distance one are shown, and the edges of $G_{2,\cup\mathbf{B}}$. The two blocks are marked in green and red.

- A set of cardinality 3 cutting off a vertex.

In the last case only the isolated vertex has a product structure (Fig. 4d).

If $n = 4$ and $d_1 = d_2 = d_3 = d_4 = 2$, then the graph G_3 is the graph of a hyper-cube. Figure 4 shows how a maximal 3-robustness structure can look like.

k -robustness implies $(k + 1)$ -robustness, and therefore $\mathcal{P}_k \subseteq \mathcal{P}_{k+1}$. This does not mean that all k -robustness structures are also $(k + 1)$ -robustness structures, for the following reason: If \mathbf{B} is a k -robustness structure and $\mathcal{S} = \cup\mathbf{B}$, then $G_{k+1,\mathcal{S}}$ may have more connected components than $G_{k,\mathcal{S}}$.

Example 28. Consider $n = 4$ binary random variables X_1, \dots, X_4 . Then

$$\mathbf{B} := \{(1, 1, 1, 1), (2, 2, 1, 1)\}, \{(1, 2, 2, 2), (2, 1, 2, 2)\}$$

is a maximal 2-robustness structure. Both elements of \mathbf{B} are connected in G_2 , but not in G_3 , see Fig. 5.

Nevertheless, the notions of l -robustness and k -robustness for $l > k$ are related as follows:

Lemma 29. Assume that $d_1 = \dots = d_n = 2$, and let \mathbf{B} be a maximal k -robustness structure of binary random variables. Then each $B \in \mathbf{B}$ is connected as a subset of G_s for all $s \leq n - 2k + 1$.

Proof. We can identify elements of \mathcal{X}_{in} with binary strings of length n . Denote by I_r the string $1 \dots 10 \dots 0$ of r ones and $n - r$ zeroes in this order. Without loss of generality assume that I_0, I_l are two elements of $B \in \mathbf{B}$, where $k \leq n - l < s \leq$

$n - 2k + 1$. Then $l \geq 2k$, and hence $\lfloor \frac{l}{2} \rfloor \geq k$. Let $m = \lceil \frac{l}{2} \rceil$. We will prove that we can replace B by $B \cup \{I_m\}$ and obtain another k -robustness structure. By maximality this will imply that I_0 and I_l are indeed connected by a path in G_s .

Otherwise there exists $A \in \mathbf{B}$, $A \neq B$, and $x \in A$ such that x and I_m agree in at least k components. Let a be the number of zeroes in the first m components of x , let b be the number of ones in the components from $m + 1$ to l and let c be the number of ones in the last $n - l$ components. Then I_m and x disagree in $a + b + c \leq n - k$ components. On the other hand, x and I_0 disagree in $(m - a) + b + c$ components, and x and I_l disagree in $a + ((l - m) - b) + c \leq a + (m - b) + c$ components. Assume that $a \geq b$ (otherwise exchange I_0 and I_l in the following argument). Then x and I_0 disagree in at most $m + c \leq \lceil \frac{l}{2} \rceil + n - l = n - \lfloor \frac{l}{2} \rfloor \leq n - k$ components, so $A \cup B$ is connected, in contradiction to the assumptions. \square

Acknowledgement. This work has been supported by the Volkswagen Foundation and the Santa Fe Institute. Nihat Ay thanks David Krakauer and Jessica Flack for many stimulating discussions on robustness.

REFERENCES

- [1] N. Ay and D. C. Krakauer, “Geometric robustness theory and biological networks,” *Theory in Biosciences*, vol. 125, no. 2, pp. 93 – 121, 2007.
- [2] G. Boldhaus, N. Bertschinger, J. Rauh, E. Olbrich, and K. Klemm, “Robustness of boolean dynamics under knockouts,” *Phys. Rev. E*, vol. 82, no. 2, p. 021916, 2010.
- [3] E. S. de Cabezón and H. Wynn, “Algebraic reliability based on monomial ideals: A review,” in *Harmony of Gröbner Bases and the Modern Industrial Society*, T. Hibi, Ed. World Scientific, 2012, pp. 314–335.
- [4] M. Drton, B. Sturmfels, and S. Sullivant, *Lectures on Algebraic Statistics*, 1st ed., ser. Oberwolfach Seminars. Birkhäuser, Basel, 2009, vol. 39.
- [5] A. Fink, “The binomial ideal of the intersection axiom for conditional probabilities,” *Journal of Algebraic Combinatorics*, vol. 33, no. 3, pp. 455–463, 2011.
- [6] J. Herzog, T. Hibi, F. Hreinsdóttir, T. Kahle, and J. Rauh, “Binomial edge ideals and conditional independence statements,” *Advances in Applied Mathematics*, vol. 45, no. 3, pp. 317 – 333, 2010.
- [7] A. S. Jarrah and R. C. Laubenbacher, “Discrete models of biochemical networks: The toric variety of nested canalizing functions,” in *Algebraic Biology*, 2007, pp. 15–22.
- [8] A. S. Jarrah, B. Raposa, and R. Laubenbacher, “Nested canalizing, unate cascade, and polynomial functions,” *Physica D*, vol. 233, no. 2, pp. 167 – 174, 2007.
- [9] S. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- [10] S. Kauffman, C. Peterson, B. Samuelsson, and C. Troein, “Genetic networks with canalizing Boolean rules are always stable,” *PNAS*, vol. 101, no. 49, pp. 17 102–17 107, 2004.
- [11] S. L. Lauritzen, *Graphical Models*, 1st ed., ser. Oxford Statistical Science Series. Oxford University Press, 1996.
- [12] M. Ohtani, “Graphs and ideals generated by some 2-minors,” *Commun. Algebra*, vol. 39, no. 3, pp. 905–917, 2011.
- [13] J. Rauh, “Generalized binomial edge ideals,” *Advances in Applied Mathematics*, accepted, 2012.
- [14] I. Swanson and A. Taylor, “Minimal primes of ideals arising from conditional independence statements,” *Preprint: arXiv:1107.5604v3*, 2011.

E-mail address: {rauh, nay}@mis.mpg.de

(J. Rauh and N. Ay) MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES, INSELSTRASSE 22, D-04103 LEIPZIG, GERMANY

(N. Ay) SANTA FE INSTITUTE, 1399 HYDE PARK ROAD, SANTA FE, NEW MEXICO 87501, USA